

Estimating Signals with Finite Rate of Innovation from Noisy Samples: A Stochastic Algorithm

Vincent Y. F. Tan, *Graduate Student Member, IEEE*, Vivek K Goyal, *Senior Member, IEEE*

Abstract

As an example of the recently-introduced concept of rate of innovation, signals that are linear combinations of a finite number of Diracs per unit time can be acquired by linear filtering followed by uniform sampling. However, in reality, samples are rarely noiseless. In this paper, we introduce a novel *stochastic* algorithm to reconstruct a signal with finite rate of innovation from its *noisy* samples. Even though variants of this problem has been approached previously, satisfactory solutions are only available for certain classes of sampling kernels, for example kernels which satisfy the Strang–Fix condition. In this paper, we consider the infinite-support Gaussian kernel, which does not satisfy the Strang–Fix condition. Other classes of kernels can be employed. Our algorithm is based on Gibbs sampling, a Markov chain Monte Carlo (MCMC) method. Extensive numerical simulations demonstrate the accuracy and robustness of our algorithm.

Index Terms

Analog-to-digital conversion, Gibbs sampling, Markov chain Monte Carlo, Sampling.

EDICS Categories: DSP-SAMP, SSP-PARE

Vincent Y. F. Tan (vtan@mit.edu) is with the Massachusetts Institute of Technology, 77 Massachusetts Avenue, Rm. 32-D570, Cambridge, MA 02139, Tel: 617-253-3816, Fax: 617-258-8364. Vincent Tan is supported by the Agency for Science, Technology and Research (A*STAR), Singapore.

Vivek K Goyal (vgoyal@mit.edu) is with the Massachusetts Institute of Technology, 77 Massachusetts Avenue, Rm. 36-690, Cambridge, MA 02139, Tel: 617-324-0367, Fax: 617-324-4290.

I. INTRODUCTION

The celebrated Nyquist-Shannon sampling theorem [1], [2]¹ states that a signal $x(t)$ known to be bandlimited to Ω_{\max} is uniquely determined by samples of $x(t)$ spaced $1/(2\Omega_{\max})$ apart. The textbook reconstruction procedure is to feed the samples as impulses to an ideal lowpass (sinc) filter. Furthermore, if $x(t)$ is not bandlimited or the samples are noisy, introducing pre-filtering by the appropriate sinc *sampling kernel* gives a procedure that finds the orthogonal projection to the space of Ω_{\max} -bandlimited signals. Thus the noisy case is handled by simple, linear, time-invariant processing.

Sampling has come a long way since the sampling theorem, but until recently the results have mostly applied only to signals contained in shift-invariant subspaces [4]. Moving out of this restrictive setting, Vetterli *et al.* [5] showed that it is possible to develop sampling schemes for certain classes of non-bandlimited signals that are not subspaces. As described in [5], for reconstruction from samples it is necessary for the class of signals to have *finite rate of innovation* (FRI). The paradigmatic example is the class of signals expressed as

$$x(t) = \sum_k c_k \phi(t - t_k)$$

where $\phi(t)$ is some known function. For each term in the sum, the signal has two real parameters c_k and t_k . If the density of t_k s (the number that appear per unit of time) is finite, the signal has FRI. It is shown constructively in [5] that the signal can be recovered from (noise-less) uniform samples of $x(t) * h(t)$ (at a sufficient rate) when $\phi(t) * h(t)$ is a sinc or Gaussian function. Results in [6] are based on similar reconstruction algorithms and greatly reduce the restrictions on the sampling kernel $h(t)$.

In practice, though, acquisition of samples is not a noiseless process. For instance, an analog-to-digital converter (ADC) has several sources of noise, including thermal noise, aperture uncertainty, comparator ambiguity, and quantization [7]. Hence, samples are inherently noisy. This motivates our central question: *Given the signal model (i.e. a signal with FRI) and the noise model, how well can we approximate the parameters that describe the signal?* In this work, we address this question and develop a novel algorithm to reconstruct the signal from the noisy samples, which we will denote $y[n]$ (see Fig. 1).

A. Related Work and Motivation

Signals with FRI were initially introduced by Vetterli *et al.* [5]. The reconstruction schemes hinged on identifying algebraically-independent parameters of the signals, *e.g.* the weights $\{c_k\}$ and time locations

¹A more expansive term could be the Whittaker-Nyquist-Kotelnikov-Shannon sampling theorem; see, e.g., [3], [4].

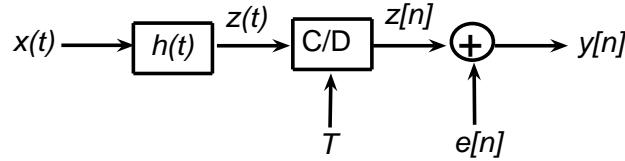


Fig. 1. Block diagram showing our problem setup. $x(t)$ is a signal with FRI given by (1) and $h(t)$ is the Gaussian filter with width σ_h given by (2). $e[n]$ is i.i.d. Gaussian noise with standard deviation σ_e and $y[n]$ are the noisy samples. From $y[n]$ we will estimate the parameters that describe $x(t)$, namely $\{c_k, t_k\}_{k=1}^K$, and σ_e , the standard deviation of the noise.

$\{t_k\}$. In the seminal paper on FRI, the sampling kernel for finite signals was chosen to be either the sinc or the Gaussian. An annihilating filter approach led to an elegant algebraic solution via polynomial root finding and least squares. The authors alluded to the noisy case and suggested the use of the Singular Value Decomposition (SVD) for dealing with noisy samples. We will show that, in fact, this method is ill-conditioned because root-finding is itself not at all robust to noise. Thus it is not amenable to practical implementations, for instance on an ADC.

Subsequently, Dragotti *et al.* [6] examined acquisition of the same signals with an eye toward implementability of the sampling kernel. Instead of using the sinc and Gaussian kernels (which do not have compact support), the authors limited the choice of kernels to functions satisfying the Strang–Fix conditions [8] (*e.g.* splines and scaling functions), exponential splines [9] and functions with rational Fourier transforms. They combined the moment-sampling and annihilating filter approaches to solve for the parameters. In our work, however, we will continue to use the Gaussian as our sampling kernel. We believe that, even though the Gaussian has infinite support, it can be well approximated by its truncated version. Hence, we can still draw insights from the analysis of using Gaussian filters and the subsequent reconstruction of the signal from its noisy samples $y[n]$. More importantly, unlike with previous approaches, the sampling kernel plays no fundamental role in the reconstruction algorithm. We use the Gaussian kernel because of its prominence in earlier work and the intuitiveness of its information spreading properties.

Maravic and Vetterli [10] and Ridolfi *et al.* [11] proposed and solved a related problem. Instead of modeling the noise at the *output*, they considered the scenario where $x(t)$, the signal in question, is corrupted by additive white noise $e(t)$. Clearly, $x_e(t) = x(t) + e(t)$ does not belong to the class of signals with FRI. However, in [10], novel algebraic/subspace-based approaches solve the sampling problem in

the Laplace domain and these methods achieve some form of optimality. In [11], various algorithms including subspace-based approaches [12] (ESPRIT and MUSIC) as well as multidimensional search methods were used and comparisons were made. The authors concluded that, in the noisy signal case, the parameters can be recovered at a rate below that prescribed by the Shannon-Nyquist Theorem but at a factor above the critical rate.

B. Our Contributions

In our paper, we solve a *different* problem. We model the noise as additive noise to the acquired samples $y[n]$, *not* the signal $x(t)$. Besides, we use the *Gaussian sampling kernel* and show that the ill-conditioning of the problem can be effectively circumvented. We demonstrate that under these conditions, we are able to estimate the parameters via a fully Bayesian approach based on Gibbs sampling (GS) [13], [14]. The prior methods are essentially *algebraic* while our algorithm is stochastic. As such, the maximization of the log-likelihood function, which we will derive in Section III, is robust to initialization.

More importantly, our algorithm is not constrained to work on the Gaussian kernel. *Any* kernel can be employed because the formulation of the Gibbs sampler does not depend on the specific form of the kernel $h(t)$. Finally, all the papers mentioned failed to estimate the standard deviation of the noise process σ_e . We address this issue in this paper.

C. Paper Organization

The rest of this paper is organized as follows: In Section II, we will formally state the problem and define the notation to be used in the rest of the paper. We proceed to delineate our algorithm: a stochastic optimization procedure based on Gibbs sampling, in Section III. We report the results of extensive numerical experiments in Section IV. In Section IV, we will also highlight some of the main deficiencies in [5], which motivate the need for new algorithms for recovering the parameters of a signal with FRI given noisy samples $y[n]$. We conclude our discussion in Section V and provide directions for further research.

II. PROBLEM DEFINITION AND NOTATION

The basic setup is shown in Fig. 1. As mentioned in the introduction, we consider a class of signals characterized by a finite number of parameters. In this paper, similar to [5], [10], [6], the class is the

weighted sum of K Diracs²

$$x(t) = \sum_{k=1}^K c_k \delta(t - t_k). \quad (1)$$

The signal to be estimated $x(t)$ is filtered using a Gaussian low-pass filter

$$h(t) = \exp\left(-\frac{t^2}{2\sigma_h^2}\right) \quad (2)$$

with width σ_h to give the signal $z(t)$. Even though $h(t)$ does not have compact support, it can be well approximated by a truncated Gaussian, which does have compact support. The filtered signal $z(t)$ is sampled at rate of $1/T$ seconds to obtain $z[n] = z(nT)$ for $n = 0, 1, \dots, N-1$. Finally, additive white Gaussian noise (AWGN) $e[n]$ is added to $z[n]$ to give $y[n]$. Therefore, the whole acquisition process from $x(t)$ to $\{y[n]\}_{n=0}^{N-1}$ can be represented by the model \mathcal{M}

$$\mathcal{M} : \quad y[n] = \sum_{k=1}^K c_k \exp\left(-\frac{(nT - t_k)^2}{2\sigma_h^2}\right) + e[n] \quad (3)$$

for $n = 0, 1, \dots, N-1$. The amount of noise added is a function of σ_e . We define the signal-to-noise ratio (SNR) in dB as

$$\text{SNR} \triangleq 10 \log_{10} \left(\frac{\sum_{n=0}^{N-1} |z[n]|^2}{\sum_{n=0}^{N-1} |z[n] - y[n]|^2} \right) \text{ dB.}$$

In the sequel, we will use boldface to denote vectors. In particular,

$$\mathbf{y} = [y[0], y[1], \dots, y[N-1]]^T, \quad (4)$$

$$\mathbf{c} = [c_1, c_2, \dots, c_K]^T, \quad (5)$$

$$\mathbf{t} = [t_1, t_2, \dots, t_K]^T. \quad (6)$$

We will sometimes use $\theta = \{\mathbf{c}, \mathbf{t}, \sigma_e\}$ to denote the complete set of decision variables. We will be measuring the performance of our reconstruction algorithms by using the normalized reconstruction error

$$\mathcal{E} \triangleq \frac{\int_{-\infty}^{\infty} |z_{est}(t) - z(t)|^2 dt}{\int_{-\infty}^{\infty} |z(t)|^2 dt}, \quad (7)$$

where $z_{est}(t)$ is the reconstructed version of $z(t)$. By construction $\mathcal{E} \geq 0$ and the closer \mathcal{E} is to 0, the better the reconstruction algorithm. In sum, the problem can be summarized as: *Given $\mathbf{y} = \{y[n] \mid n = 0, \dots, N-1\}$ and the model \mathcal{M} , estimate the parameters $\{c_k, t_k\}_{k=1}^K$ to minimize \mathcal{E} . Also estimate the noise variance σ_e^2 .*

²The use of a Dirac delta simplifies the discussion. It can be replaced by a known pulse $g(t)$ and then absorbed into the sampling kernel $h(t)$, yielding an effective sampling kernel $g(t) * h(t)$.

III. PRESENTATION OF THE GIBBS SAMPLER

In this section, we will describe the stochastic optimization procedure based on Gibbs sampling to estimate $\theta = \{\mathbf{c}, \mathbf{t}, \sigma_e\}$.

A. Gibbs Sampling (GS)

Markov chain Monte Carlo (MCMC) in the form of the Gibbs sampler, and the Metropolis-Hastings algorithm allows any distribution to be simulated on a finite-dimensional state space specified by any conditional density. The Gibbs sampler was first studied by the statistical physics community [15] and then later in the statistics community [13], [16], [17]. The basis for Gibbs sampling is the Hammersley-Clifford theorem [18] which states that given the data \mathbf{y} , the conditional densities $p_i(\theta_i | \theta_{\{j \neq i\}}, \mathbf{y}, \mathcal{M})$ contain sufficient information to produce samples from the joint density $p(\theta | \mathbf{y}, \mathcal{M})$. Furthermore, the joint density can be directly derived from the conditional densities.

Gibbs sampling has been used extensively and successfully in image [13] and audio restoration [14]. The Gibbs sampler is presented here to estimate $\theta = \{\mathbf{c}, \mathbf{t}, \sigma_e\}$. To simulate our Gibbs sampler, we use the i.i.d. Gaussian noise assumption and the model in (3) to express the log-likelihood of the parameters given the observations as:

$$\begin{aligned} & \log p(\mathbf{c}, \mathbf{t}, \sigma_e | \mathbf{y}, \mathcal{M}) \\ & \propto -(N+1) \log(\sigma_e) \\ & \quad - \frac{1}{2\sigma_e^2} \sum_{n=0}^{N-1} \left[y[n] - \sum_{k=1}^K c_k \exp\left(-\frac{(nT-t_k)^2}{2\sigma_h^2}\right) \right]^2. \end{aligned} \quad (8)$$

A Jeffrey's (improper) non-informative prior has been assigned to the standard deviation of the noise such that

$$p(\sigma_e) \propto \frac{1}{\sigma_e}. \quad (9)$$

In the Gibbs sampling algorithm, as soon as a variate is drawn, it is inserted immediately into the conditional p.d.f. and it remains there until being substituted in the next iteration. This is shown in the following algorithm.³

Require: $\mathbf{y}, I, I_b, \theta^{(0)} = \{\mathbf{c}^{(0)}, \mathbf{t}^{(0)}, \sigma_e^{(0)}\}$

for $i \leftarrow 1 : I + I_b$ **do**

³For brevity, the dependence on the model \mathcal{M} is omitted from the conditional density expressions.

$$\begin{aligned}
c_1^{(i)} &\sim p(c_1|c_2^{(i-1)}, c_3^{(i-1)}, \dots, c_K^{(i-1)}, \mathbf{t}^{(i-1)}\sigma_e^{(i-1)}, \mathbf{y}) \\
c_2^{(i)} &\sim p(c_2|c_1^{(i)}, c_3^{(i-1)}, \dots, c_K^{(i-1)}, \mathbf{t}^{(i-1)}\sigma_e^{(i-1)}, \mathbf{y}) \\
&\vdots \sim \vdots \\
c_K^{(i)} &\sim p(c_K|c_1^{(i)}, c_2^{(i)}, \dots, c_{K-1}^{(i)}, \mathbf{t}^{(i-1)}, \sigma_e^{(i-1)}, \mathbf{y}) \\
t_1^{(i)} &\sim p(t_1|\mathbf{c}^{(i)}, t_2^{(i-1)}, t_3^{(i-1)}, \dots, t_K^{(i-1)}, \sigma_e^{(i-1)}, \mathbf{y}) \\
t_2^{(i)} &\sim p(t_2|\mathbf{c}^{(i)}, t_1^{(i)}, t_3^{(i-1)}, \dots, t_K^{(i-1)}, \sigma_e^{(i-1)}, \mathbf{y}) \\
&\vdots \sim \vdots \\
t_K^{(i)} &\sim p(t_K|\mathbf{c}^{(i)}, t_1^{(i)}, t_2^{(i)}, \dots, t_{K-1}^{(i)}, \sigma_e^{(i-1)}, \mathbf{y}) \\
\sigma_e^{(i)} &\sim p(\sigma_e|\mathbf{c}^{(i)}, \mathbf{t}^{(i)}, \mathbf{y})
\end{aligned}$$

end for

Compute $\hat{\boldsymbol{\theta}}_{\text{MMSE}}$ using (10)

return $\hat{\boldsymbol{\theta}}_{\text{MMSE}}$

In the algorithm, $\vartheta \sim \bar{p}(\cdot)$ means that ϑ is a random draw from $\bar{p}(\cdot)$. The superscript number (i) denotes the current iteration. After I_b iterations⁴ the Markov chain approximately reaches its stationary distribution $p(\boldsymbol{\theta}|\mathbf{y}, \mathcal{M})$. Minimum mean squared error (MMSE) estimates can then be approximated by taking averages of the samples from the next I iterations $\{\boldsymbol{\theta}^{(I_b+1)}, \boldsymbol{\theta}^{(I_b+2)}, \dots, \boldsymbol{\theta}^{(I_b+I)}\}$, *i.e.*,

$$\hat{\boldsymbol{\theta}}_{\text{MMSE}} = \int \boldsymbol{\theta} p(\boldsymbol{\theta}|\mathbf{y}, \mathcal{M}) d\boldsymbol{\theta} \approx \frac{1}{I} \sum_{i=I_b+1}^{I_b+I} \boldsymbol{\theta}^{(i)}. \quad (10)$$

B. Presentation of the Posterior Densities in the GS

We will now derive the conditional densities. In the sequel, we will use the notation $\boldsymbol{\theta}_{-\ell}$ to denote the set of parameters excluding the ℓ th parameter. It follows from Bayes' theorem that

$$p(\theta_\ell|\boldsymbol{\theta}_{-\ell}, \mathbf{y}, \mathcal{M}) \propto p(\mathbf{y}|\boldsymbol{\theta}, \mathcal{M}) p(\boldsymbol{\theta}). \quad (11)$$

Thus, the required conditional distributions are proportional to the *likelihood* of the data times the *priors* on the parameters. The likelihood function of \mathbf{y} given the model is given in (8) from the Gaussian noise assumption. Thus, we can calculate the posterior distributions of the parameters given the rest of the parameters. The parameters conditioned on are taken as constant and can be left out of the posterior. We will sample from these posterior densities in the GS iterations as shown in the above algorithm.

We will now proceed to present the posterior densities. The derivations are provided in the Appendix.

⁴ I_b is also commonly known as the *burn-in* period in the Gibbs sampling and MCMC literature [14].

1) *Sampling c_k :* c_k is sampled from a Gaussian distribution given by

$$p(c_k | \boldsymbol{\theta}_{-c_k}, \mathbf{y}, \mathcal{M}) = \mathcal{N} \left(c_k; -\frac{\beta_k}{2\alpha_k}, \frac{1}{2\alpha_k} \right), \quad (12)$$

where

$$\alpha_k \triangleq \frac{1}{2\sigma_e^2} \sum_{n=0}^{N-1} \exp \left(-\frac{(nT - t_k)^2}{\sigma_h^2} \right), \quad (13)$$

$$\begin{aligned} \beta_k \triangleq & \frac{1}{\sigma_e^2} \sum_{n=0}^{N-1} \exp \left(-\frac{(nT - t_k)^2}{2\sigma_h^2} \right) \\ & \cdot \left\{ \sum_{\substack{k'=1 \\ k' \neq k}}^K c_{k'} \exp \left(-\frac{(nT - t_{k'})^2}{2\sigma_h^2} \right) - y[n] \right\}. \end{aligned} \quad (14)$$

It is easy to sample from Gaussian densities when the parameters (α_k, β_k) have been determined.

2) *Sampling t_k :* t_k is sampled from a distribution of the form

$$\begin{aligned} p(t_k | \boldsymbol{\theta}_{-t_k}, \mathbf{y}, \mathcal{M}) & \propto \exp \left[-\frac{1}{2\sigma_e^2} \sum_{n=0}^{N-1} \gamma_k \exp \left(-\frac{(nT - t_k)^2}{\sigma_h^2} \right) \right. \\ & \quad \left. + \nu_k \exp \left(-\frac{(nT - t_k)^2}{2\sigma_h^2} \right) \right] \end{aligned} \quad (15)$$

where

$$\gamma_k \triangleq c_k^2, \quad (16)$$

$$\nu_k \triangleq 2c_k \left\{ \sum_{\substack{k'=1 \\ k' \neq k}}^K c_{k'} \exp \left(-\frac{(nT - t_{k'})^2}{2\sigma_h^2} \right) - y[n] \right\}. \quad (17)$$

It is not straightforward to sample from this distribution. We can sample t_k from a uniform grid of discrete values with probability masses proportional to (15). But in practice, and for greater accuracy, we used rejection sampling [19], [20] to generate samples $t_k^{(i)}$ from $p(t_k | \boldsymbol{\theta}_{-t_k}, \mathbf{y}, \mathcal{M})$. The proposal distribution $\tilde{q}(t_k)$ was chosen to be an appropriately scaled Gaussian, since it is easy to sample from Gaussians. This is shown in the following algorithm.

Require: $\tilde{p}(t_k) \triangleq p(t_k | \boldsymbol{\theta}_{-t_k}, \mathbf{y}, \mathcal{M})$

Select $\tilde{q}(t_k) \sim \mathcal{N}$ and c s.t. $\tilde{p}(t_k) < c\tilde{q}(t_k)$

$u \sim \mathcal{U}(0, 1)$

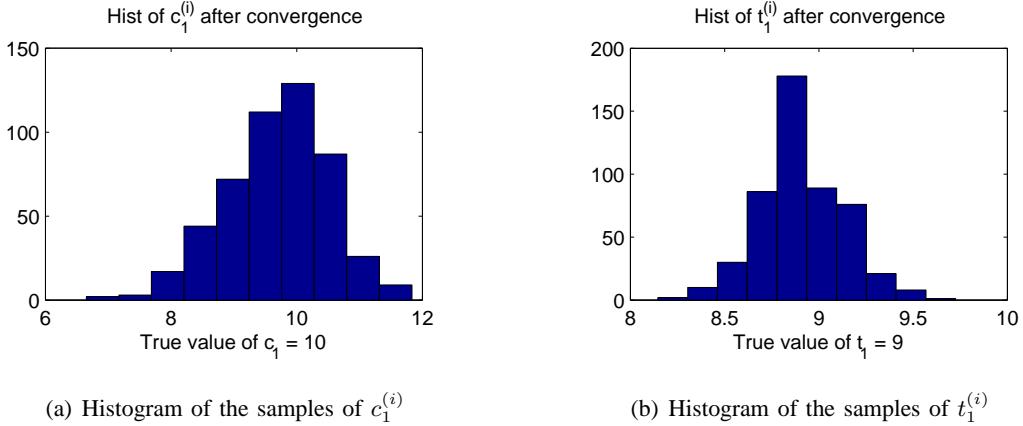


Fig. 2. Note that the variance of the stationary distribution of the t_k s is smaller than that of the c_k s after convergence of the Markov chain.

repeat

$$t_k \sim \tilde{q}(t_k)$$

until $u < \tilde{p}(t_k)/(c\tilde{q}(t_k))$

3) *Sampling σ_e :* σ_e is sampled from the ‘Square-root Inverted-Gamma’ [21] distribution $\mathcal{IG}^{-1/2}(\sigma_e; \varphi, \lambda)^5$,

$$p(\sigma_e | \boldsymbol{\theta}_{-\sigma_e}, \mathbf{y}, \mathcal{M}) = \frac{2\lambda^\varphi \sigma_e^{-(2\varphi+1)}}{\Gamma(\varphi)} \exp\left(-\frac{\lambda}{\sigma_e^2}\right) \mathbb{I}_{[0,+\infty)}(\sigma_e), \quad (18)$$

where

$$\varphi \triangleq \frac{N}{2}, \quad (19)$$

$$\lambda \triangleq \frac{1}{2} \left[y[n] - \sum_{k=1}^K c_k \exp\left(-\frac{(nT - t_k)^2}{2\sigma_h^2}\right) \right]^2. \quad (20)$$

Thus the distribution of the variance of the noise σ_e^2 is Inverted Gamma, which corresponds to the conjugate prior of σ_e^2 in the expression of $\mathcal{N}(e; 0, \sigma_e^2)$ [21] and thus it is easy to sample from. In our simulations, we sampled from this density using the Matlab function `gamrnd` and applied the ‘Inverted Square-root’ transformation

C. Further Improvements via Linear Least Squares Estimation

We can perform an additional post-processing step to improve on the estimates of c_k . We noted from our preliminary experiments (see Fig. 2) that the variance of the stationary distribution of the t_k s is smaller

⁵ X follows a ‘Square-root Inverted-Gamma’ distribution if X^{-2} follows a Gamma distribution.

Param.	K	σ_e	N	SNR
Value	5	0 and 10^{-6}	30	∞ and 137 dB

TABLE I

PARAMETER VALUES FOR DEMONSTRATION OF THE ANNIHILATING FILTER AND ROOT-FINDING ALGORITHM.

than that of the c_k s. This results in better estimates for the locations t_k s as compared to the magnitudes c_k s. Now, we observe that $y[n]$, the observations, are linear in the c_k s once the t_k s are known. A natural extension to our GS algorithm is to augment our c_k estimates with a linear least squares estimation (LLSE) procedure using \mathbf{y} and the MMSE estimates of t_k . Eqn. (3) can be written as

$$y[n] = \sum_{k=1}^K c_k h(nT - t_k) + e[n], \quad 0 \leq n \leq N - 1 \quad (21)$$

with $h(t)$, the Gaussian sampling kernel, given in (2). Given the set of estimates of the time locations $\{\hat{t}_k\}_{k=1}^K$, we can rewrite (21) as a matrix equation, giving

$$\mathbf{y} = \mathbf{H}\mathbf{c} + \mathbf{e},$$

where $[\mathbf{H}]_{nk} = h(nT - \hat{t}_k)$ and $1 \leq n \leq N$, $1 \leq k \leq K$. We now minimize the square of the residual $\|\mathbf{e}\|^2 = \|\mathbf{H}\mathbf{c} - \mathbf{y}\|^2$, giving the normal equations $\mathbf{H}^T \mathbf{H} \mathbf{c} = \mathbf{H}^T \mathbf{y}$ and the least squares solution [22]

$$\hat{\mathbf{c}}_{\text{LS}} = (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \mathbf{y}. \quad (22)$$

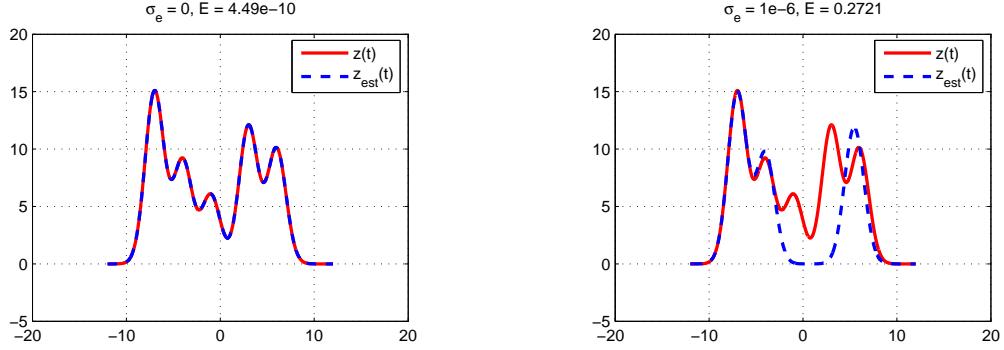
From our experiments, we found that, in general, using $\hat{\mathbf{c}}_{\text{LS}}$ as estimates for the magnitudes of the impulses provided a lower reconstruction error \mathcal{E} .

IV. NUMERICAL RESULTS AND EXPERIMENTS

In this section, we will first review the annihilating filter and root-finding method algorithm for solving for the parameters of a signal with FRI. This algorithm provides a baseline for comparison. Then we will provide extensive simulation results to validate the accuracy of the algorithm we proposed in Section III.

A. Problems with Annihilating Filter and Root-Finding

In [5], Vetterli *et al.* introduced the concept of a class of signals with a finite rate of innovation. For signals of the form (1) and certain sampling kernels, the annihilating filter was used as a means to locate the t_k values. Subsequently a least squares approach yielded the weights c_k . It was shown that in the



(a) The annihilating filter approach reconstructs the signal exactly in the noiseless scenario.
(b) The reconstruction completely breaks down when noise of a small standard deviation $\sigma_e = 10^{-6}$ (SNR=137 dB) is added.

Fig. 3. Demonstration of the annihilating filter/root-finding approach.

noiseless scenario, this method recovers the parameters exactly (see Fig. 3(a)). For completeness, we will briefly outline their method here. Denoting the noiseless samples by $z[n]$, (3) can be written as

$$p[n] = \sum_{k=1}^K a_k u_k^n, \quad n = 0, 1, \dots, N-1, \quad (23)$$

with the identifications

$$p[n] = \exp(n^2 T^2 / (2\sigma_h^2)) z[n], \quad (24)$$

$$a_k = c_k \exp(-t_k^2 / (2\sigma_h^2)), \quad (25)$$

$$u_k = \exp(t_k T / \sigma_h^2). \quad (26)$$

Now, since $p[n]$ is a linear combination of exponentials, we find the annihilating filter $a[n]$ such that

$$a[n] * p[n] = \sum_{\ell=0}^K a[\ell] p[n-\ell] = 0, \quad \forall n \in \mathbb{Z}.$$

This can be written in matrix/vector form as $\mathbf{P}\mathbf{a} = \mathbf{0}$. This system will admit a solution when $\text{rank}(\mathbf{P}) = K$. In practice this is solved using an SVD where $\mathbf{P} = \mathbf{U}\Sigma\mathbf{V}^T$ and $\mathbf{a} = \mathbf{V}\mathbf{e}_{K+1}$ and \mathbf{e}_{K+1} is a length- $(K+1)$ vector with 1 in position $(K+1)$ and 0 elsewhere. Now, once the coefficients $a[n]$ are found, the values u_k are simply the roots of the filter

$$A(z) = \sum_{n=0}^K a[n] z^{-n}.$$

The t_k s can then be determined from (25) and the solution for the c_k s essentially parallels the development in Section III-C.

In the same paper, it was suggested that to deal with the noisy samples, we can minimize $\|\mathbf{P}\mathbf{a}\|$, in which case, \mathbf{a} is the eigenvector that corresponds to the smallest eigenvalue of $\mathbf{P}^T\mathbf{P}$. Here, we argue that this method is inherently ill-conditioned and thus not robust to noise.

- 1) Firstly, minimizing $\|\mathbf{P}\mathbf{a}\|$ involves finding the eigenvector \mathbf{v}_1 that corresponds to the largest eigenvalue λ_1 . Because computing eigenvalues and eigenvectors are essentially root-finding operations, this is ill-conditioned.
- 2) Secondly, even if the vector $\mathbf{a} = \mathbf{v}_1$ can be found, the zeros of the filter $A(z)$ have to be found. This again involves root finding, which is ill-conditioned.
- 3) Finally, from (24), any noise added to $z[n]$ will be exponentially weighted in the observations $p[n]$. We feel that this is the greatest source of ill-conditioning.

Because of the three reasons highlighted above, there is a need to explore new algorithms for finding the parameters. In Fig. 3, we show a simulation with the parameters as tabulated in Table I, but we varied the noise ($\sigma_e = 10^{-6}$ gives SNR = 137 dB, a very low noise level). We observe from Fig. 3(b) that (without oversampling) the annihilating filter and root-finding method is not robust even when a minuscule amount of noise is added.

Remark The root-finding method is so unstable that, at times, even for low levels of noise, we obtain complex roots for the locations $\{t_k\}_{k=0}^{K-1}$. To solve this problem, we orthogonally projected the polynomial described by the filter coefficients $a[n]$ to the closest polynomial that belongs to the space of polynomials with real roots only.

B. Performance of our GS Algorithm

Clearly, the annihilating filter/root-finding algorithm is not robust to noise. We have suggested an alternative reconstruction algorithm in Section III, and in this section, we will present our results on several synthetic examples.⁶

1) Initial Demonstration: To demonstrate the evolution the Gibbs sampler, we performed an initial experiment and chose the parameters to be those in Table I, with the exception that the noise standard deviation was increased to $\sigma_e = 2.5$, giving an SNR of 10.2 dB. We plot the iterates in Fig. 4. The true filtered signal $z(t)$ and its estimate $z_{est}(t)$ are plotted in Fig. 5. We note the close similarity between $z(t)$ and $z_{est}(t)$.

⁶All the code, written in MATLAB, can be found at the first author's homepage <http://web.mit.edu/vtan/frimcmc>.

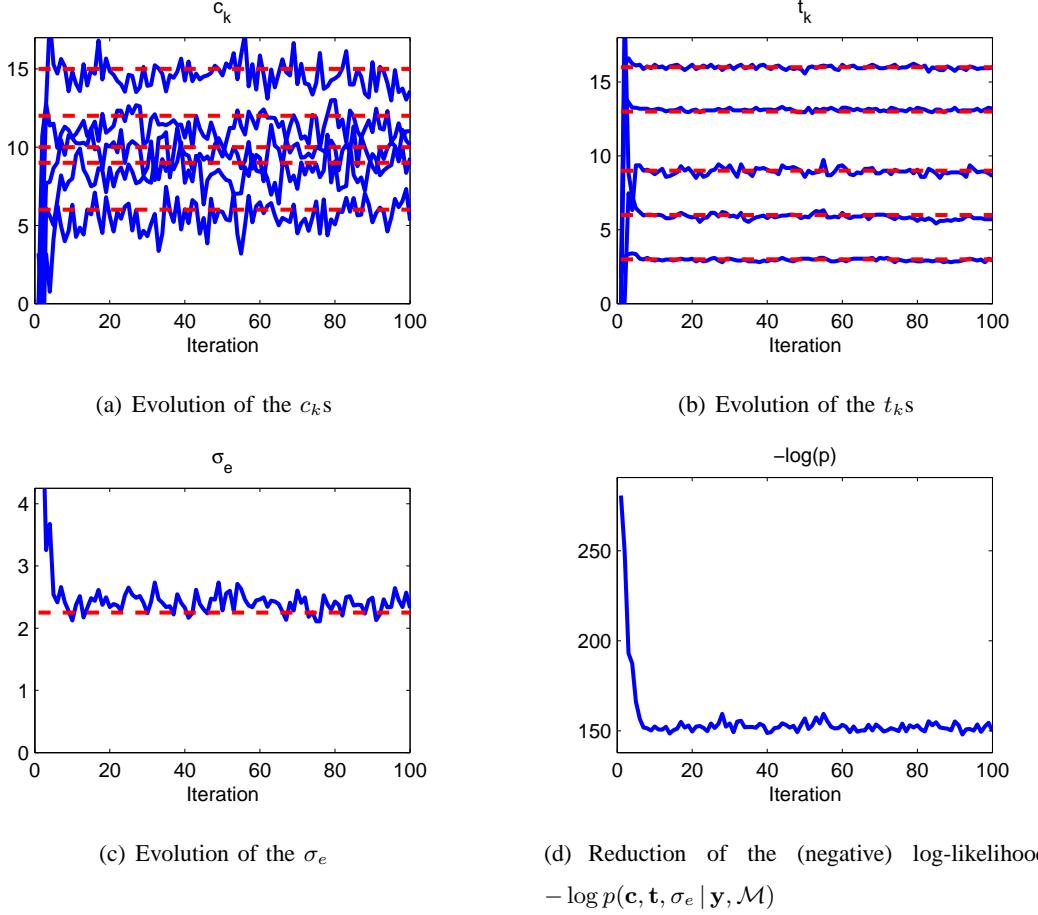


Fig. 4. Evolution of the GS algorithm. The iterates of the parameters $\{c_k, t_k\}_{k=1}^K$ and σ_e are shown. The true values are indicated by the broken red lines. In Fig. 4(d), we see that the negative log-likelihood converges to the global minimum in fewer than 20 iterations for this problem size ($K = 5$).

We observe that the sampler converges in fewer than 20 iterations for this run, even though the parameter values were initialized far from their optimal values. We emphasize that as GS is essentially a stochastic optimization procedure (not unlike Simulated Annealing or Genetic Algorithms), it is insensitive to the choice of starting point $\theta^{(0)}$. The Markov Chain is guaranteed to converge to the stationary distribution after the burn-in period [19].

2) *Further Experiments:* To further validate our algorithm, we performed extensive simulations (Expts A and B) on two different problem sizes to validate our algorithm. For consistency, each experiment was repeated using 100 different random seeds and the means of \mathcal{E} [cf. (7)] taken. The parameters are chosen according to Table II. The unknown parameters were initialized as $\mathbf{c}^{(0)} = \mathbf{t}^{(0)} = [0, \dots, 0]$ and $\sigma_e^{(0)} = 0.01$. The results for Expts A and B are shown in Fig. 7(a) and 7(b) respectively. We noted from

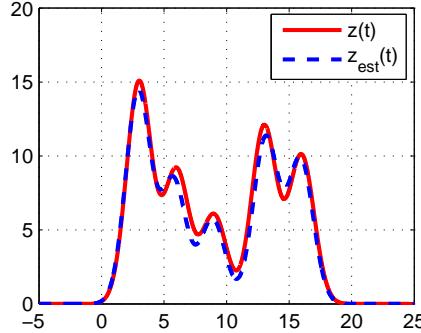


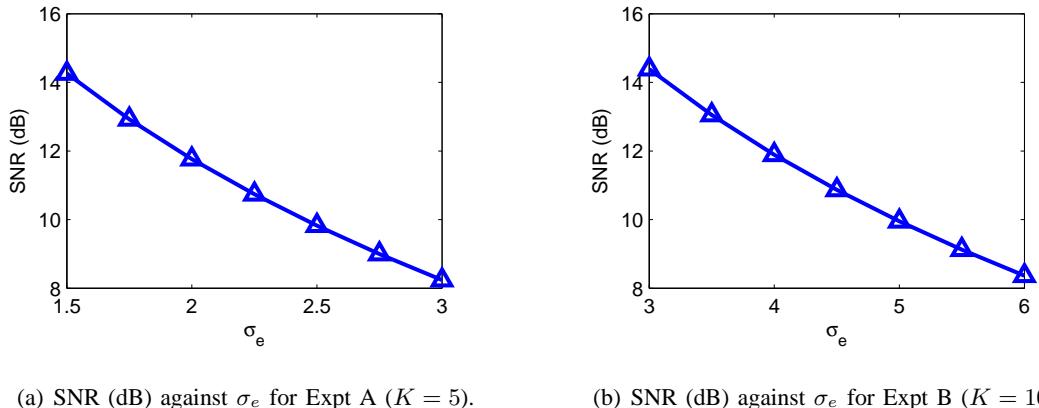
Fig. 5. Comparison between $z(t)$ and $z_{est}(t)$ using the GS algorithm. For this run, $\mathcal{E} = 0.0072$.

Param.	K	σ_e	N	SNR	\mathcal{E}
Expt A	5	1.5:0.25:3.0	50:25:150	Fig 6(a)	Fig 7(a)
Expt B	10	3.0:0.50:6.0	100:50:250	Fig 6(b)	Fig 7(b)

TABLE II
PARAMETER VALUES FOR NUMERICAL SIMULATIONS.

these experiments that:

- The GS algorithm is insensitive to initialization. It *always* finds approximately optimal estimates from any starting point because the Markov chain provably converges to the stationary distribution [19].
- The LLSE post-processing step in the GS algorithm reduces the reconstruction error \mathcal{E} . This is a



(a) SNR (dB) against σ_e for Expt A ($K = 5$).

(b) SNR (dB) against σ_e for Expt B ($K = 10$).

Fig. 6. SNR (dB) against σ_e for the two experiments.

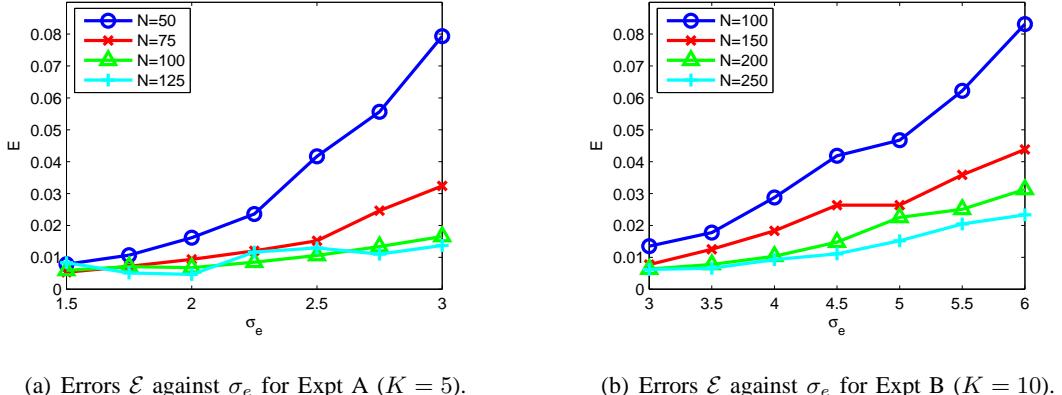


Fig. 7. Plots of E against σ_e for various oversampling factors and problem sizes.

consequence of using the (more accurate) t_k s from the sampler to estimate the c_k s via LLSE, instead of using the c_k s from the sampler directly.

- From the two plots in Fig. 7, we observe that, if the problem size doubles (from $K = 5$ to $K = 10$), with corresponding doubling of (σ_e, N) , E remains approximately constant. This insures scalability of the algorithm. For example, $E(K = 5, \sigma_e = 2.5, N = 50) \approx E(K = 10, \sigma_e = 5.0, N = 100) \approx 0.045$.
- The noise standard deviation σ_e can be estimated accurately in the GS algorithm as shown in Fig. 4(c). This may be important in some applications.

To conclude, even though the annihilating filter approach [5] is more computationally efficient than our algorithms, it is certainly not amenable to scenario where noisy samples are acquired.

V. CONCLUSIONS

In this paper, we addressed the problem of reconstructing a signal with FRI given noisy samples. We showed that it is possible to circumvent some of the problems of the annihilating filter and root-finding approach [5]. We introduced the Gibbs sampling algorithm. From the performance plots, we observe that GS performs very well as compared to the annihilating filter method, which is not robust to noise.

Perhaps the most important observation we made is the following: The success of the *fully Bayesian* GS algorithm does not depend on the choice of kernel $h(t)$. The formulation of the GS does not depend on the specific form of $h(t)$. In fact, we used a Gaussian sampling kernel to illustrate that our algorithm is not restricted to the classes of kernels considered in [6].

A natural extension to our work here is to assign structured priors to \mathbf{c} , \mathbf{t} and σ_e . These priors can

themselves be dependent on their own set of *hyperparameters*, giving a hierarchical Bayesian formulation. In this way, there would be greater flexibility in the parameter estimation process. We can also seek to improve on the computational load of the algorithms introduced here. Another interesting research direction is to examine the feasibility of using the subspace-based approaches [10] to solve the problem of acquired samples that are noisy.

A question that remains is: *How well can real-world signals (including natural images) be modeled as signals with FRI?* We believe the answer will have profound ramifications for areas such as sparse approximation [23] and compressed sensing [24], [25].

APPENDIX

DERIVATION OF THE CONDITIONAL DENSITIES

For brevity, we define

$$g_{nk} \triangleq h(nT - t_k) = \exp \left(-\frac{(nT - t_k)^2}{2\sigma_h^2} \right).$$

We start from the log-likelihood of the parameters $\boldsymbol{\theta}$ given the data \mathbf{y} and model \mathcal{M} [cf. (8)]. To obtain $p(c_k|\boldsymbol{\theta}_{-c_k}, \mathbf{y}, \mathcal{M})$, we treat the other parameters $\boldsymbol{\theta}_{-c_k}$ as constant, giving $\log p(c_k|\boldsymbol{\theta}_{-c_k}, \mathbf{y}, \mathcal{M})$ proportional to

$$-\frac{1}{2\sigma_e^2} \sum_{n=0}^{N-1} \left[c_k^2 g_{nk}^2 + 2c_k g_{nk} \left(\sum_{\substack{k'=1 \\ k' \neq k}}^K c_{k'} g_{nk'} - y[n] \right) \right].$$

Comparing this expression in c_k to the Gaussian distribution with mean μ and variance σ^2 ,

$$\log p(c_k; \mu, \sigma^2) \propto -\frac{1}{2\sigma^2} (c_k - \mu)^2,$$

and equating coefficients, we obtain (13) and (14). The distribution $p(t_k|\boldsymbol{\theta}_{-t_k}, \mathbf{y}, \mathcal{M})$ can be obtained similarly and is omitted. Finally for the noise standard deviation σ_e ,

$$\log p(\sigma_e|\boldsymbol{\theta}_{-\sigma_e}, \mathbf{y}, \mathcal{M}) \propto -(N+1) \log(\sigma_e) - \frac{\lambda}{\sigma^2},$$

where λ is defined in (20). Taking the antilog on both sides yields

$$p(\sigma_e|\boldsymbol{\theta}_{-\sigma_e}, \mathbf{y}, \mathcal{M}) \propto \sigma_e^{-(N+1)} \exp \left(-\frac{\lambda}{\sigma^2} \right),$$

which is the ‘Square-root Inverted-Gamma’ distribution with parameters given by (19) and (20). All the densities have been derived.

REFERENCES

- [1] C. E. Shannon, "Communication in the presence of noise," *Proc. Institute of Radio Engineers*, vol. 37, no. 1, pp. 10–21, Jan. 1949.
- [2] H. Nyquist, "Certain topics in telegraph transmission theory," *Trans. American Institute of Electrical Engineers*, vol. 47, pp. 617–644, Apr. 1928.
- [3] A. J. Jerri, "The Shannon sampling theorem—its various extensions and applications: A tutorial review," *Proc. IEEE*, vol. 65, pp. 1565–1596, Nov. 1977.
- [4] M. Unser, "Sampling—50 years after Shannon," *Proc. IEEE*, vol. 88, no. 4, pp. 569–587, 2000.
- [5] M. Vetterli, P. Marziliano, and T. Blu, "Sampling signals with finite rate of innovation," *IEEE Trans. Signal Processing*, vol. 50, no. 6, pp. 1417–1428, 2002.
- [6] P. L. Dragotti, M. Vetterli, and T. Blu, "Sampling moments and reconstructing signals of finite rate of innovation: Shannon meets Strang–Fix," *IEEE Trans. Signal Processing*, vol. 55, no. 5, pp. 1741–1757, 2007.
- [7] R. H. Walden, "Analog-to-digital converter survey and analysis," *IEEE J. Selected Areas of Communication*, vol. 17, no. 4, pp. 539–550, Apr. 1999.
- [8] G. Strang and G. Fix, "A Fourier analysis of the finite element variational method," *Constructive Aspects of Functional Analysis*, pp. 796–830, 1971.
- [9] M. Unser and T. Blu, "Cardinal exponential splines: Part I – Theory and filtering algorithms," *IEEE Trans. Signal Processing*, vol. 53, no. 4, pp. 1425–1438, Apr. 2005.
- [10] I. Maravic and M. Vetterli, "Sampling and reconstruction of signals with finite rate of innovation in the presence of noise," *IEEE Trans. Signal Processing*, vol. 53, no. 8, pp. 2788–2805, 2005.
- [11] A. Ridolfi, I. Maravic, J. Kusuma, and M. Vetterli, "Sampling signals with finite rate of innovation: The noisy case," *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2002.
- [12] P. Stoica and R. Moses, *Introduction to Spectral Analysis*. Prentice Hall, 1997.
- [13] S. Geman and D. Geman, "Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 6, pp. 721–741, 1984.
- [14] S. J. Godsill and P. J. W. Rayner, *Digital Audio Restoration: A Statistical Model Based Approach*. Springer-Verlag, 1998.
- [15] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and W. Teller, "Equations of state calculations by fast computing machines," *J. Chemical Physics*, vol. 21, pp. 1087–1091, 1953.
- [16] W. K. Hastings, "Monte Carlo sampling methods using Markov chains and their applications," *Biometrika*, vol. 57, pp. 97–109, 1970.
- [17] W. J. Fitzgerald, S. J. Godsill, A. C. Kokaram, and J. A. Stark, *Bayesian Methods in Signal and Image Processing*. Oxford Univ. Press, 1999, ch. Bayesian Statistics.
- [18] J. M. Hammersley and M. S. Clifford, "Markov fields on finite graphs and lattices," *Unpublished*, 1970.
- [19] L. Tierney, "Markov chains for exploring posterior distributions," School of Statistics, Univ. of Minnesota, Tech. Rep. 560, Mar. 1994.
- [20] C. P. Robert and G. Casella, *Monte Carlo Statistical Methods*, 2nd ed. New York: Springer-Verlag, 2004.
- [21] J. M. Bernardo and A. F. M. Smith, *Bayesian Theory*, 1st ed. Wiley, 2001.
- [22] G. Strang, *Introduction to Linear Algebra*, 3rd ed. Wellesley Cambridge Press, 2001.
- [23] D. L. Donoho, M. Vetterli, R. A. DeVore, and I. Daubechies, "Data compression and harmonic analysis," *IEEE Trans. Information Theory*, vol. 44, no. 6, pp. 2435–2476, Oct. 1998.

- [24] D. L. Donoho, "Compressed sensing," *IEEE Trans. Information Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.
- [25] E. J. Candès and T. Tao, "Near-optimal signal recovery from random projections: Universal encoding strategies?" *IEEE Trans. Information Theory*, vol. 52, no. 12, pp. 5406–5425, Dec. 2006.